RUTA A TRAVÉS DE LA CIENCIA DE DATOS

Deep Learning en 25 artículos for the working scientist

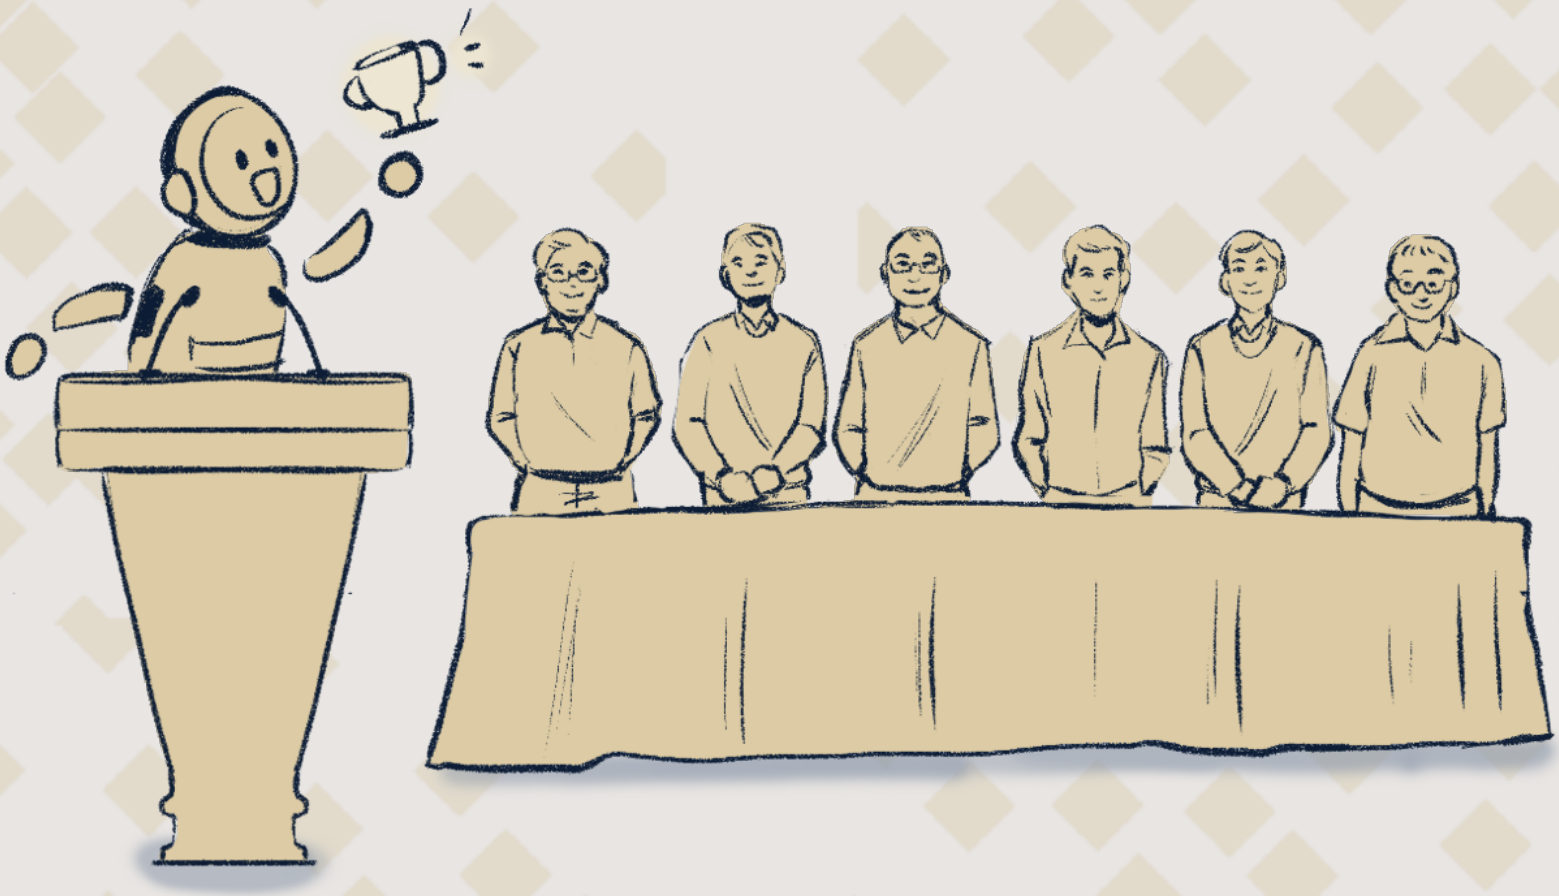

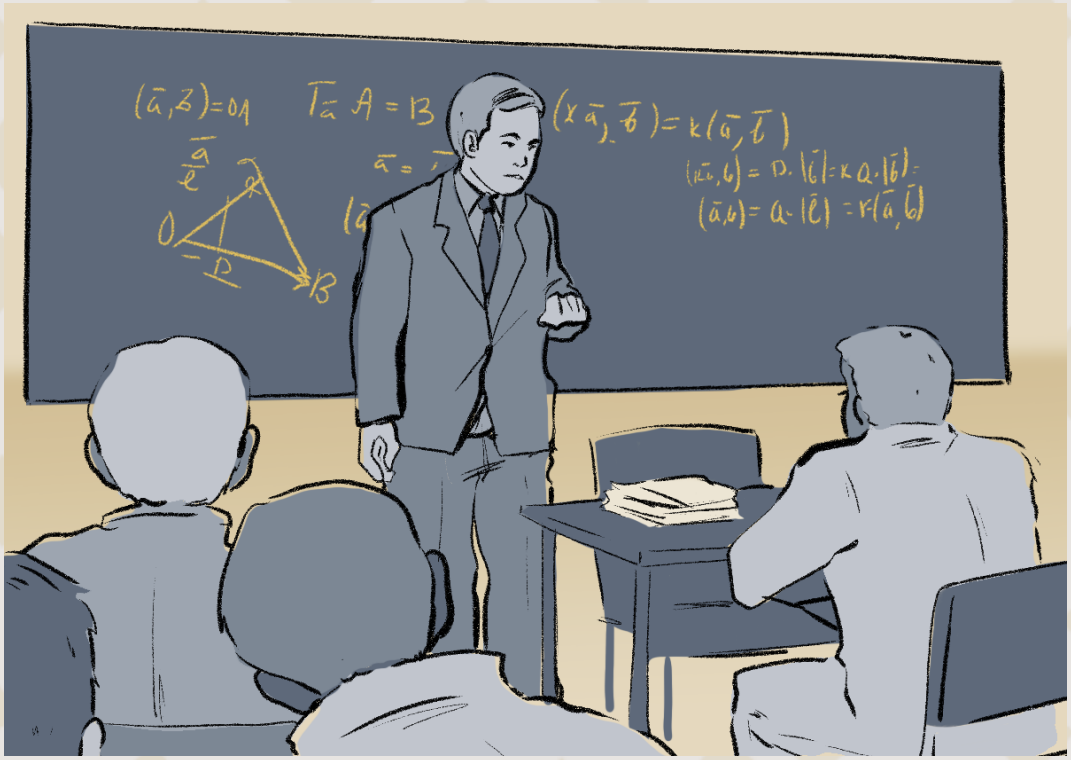

Esta curaduría de **25 artículos** ha sido diseñada como una guía de lectura paralela al curso *Deep Learning (Avanzado)*, con el objetivo de fortalecer la comprensión técnica y práctica de las redes neuronales profundas y sus aplicaciones más importantes en la industria.

En esta presentación **nos hemos centrado exclusivamente en artículos que tratan arquitecturas, algoritmos y técnicas específicas del Deep Learning moderno**, omitiendo aquellos trabajos de estadística clásica o de machine learning tradicional que abordaremos en otras selecciones temáticas.

Asimismo, hemos dejado fuera los textos cuyo enfoque es principalmente matemático-formal —como los dedicados a pruebas de convergencia, complejidad computacional o marcos bayesianos rigurosos—, los cuales formarán parte de nuestra colección *Deep Learning en 25 artículos for the Working Mathematician*.

1.

Deep learning. Nature, 521(7553), pp.436–444.

2.

Generative adversarial nets. NeurIPS.

3.

Attention is All You Need. NeurIPS.

4.

Sequence to sequence learning with neural networks. NeurIPS.

5.

Neural machine translation by jointly learning to align and translate. ICLR.

6.

Reducing the dimensionality of data with neural networks. Science, 313(5786), pp.504–507.

7.

Improving language understanding by generative pre-training.

8.

Language Models are Few-Shot Learners. NeurIPS.

9.

BERT: Pre-training of deep bidirectional transformers for language understanding. NAACL.

10.

An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. ICLR.

11.

Estimating Counterfactual Treatment Outcomes over Time Through Adversarially Balanced Representations. ICML.

12.

Machine Learning. Springer.

13.

Google's Neural Machine Translation System: Bridging the Gap between Human and Machine Translation.

14.

Effective approaches to attention-based neural machine translation. EMNLP.

15.

Continuous control with deep reinforcement learning. arXiv preprint.

16.

Proximal Policy Optimization Algorithms. arXiv preprint.

17.

Semi-Supervised Classification with Graph Convolutional Networks. ICLR.

18.

Inductive Representation Learning on Large Graphs. NeurIPS.

19.

FastGCN: Fast learning with graph convolutional networks via importance sampling. ICLR.

20.

Image style transfer using convolutional neural networks. CVPR.

21.

Deep residual learning for image recognition. CVPR.

22.

Very deep convolutional networks for large-scale image recognition. ICLR.

23.

Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. ICML.

24.

Auto-Encoding Variational Bayes. ICLR.

25.

Gradient-based learning applied to document recognition. Proceedings of the IEEE.

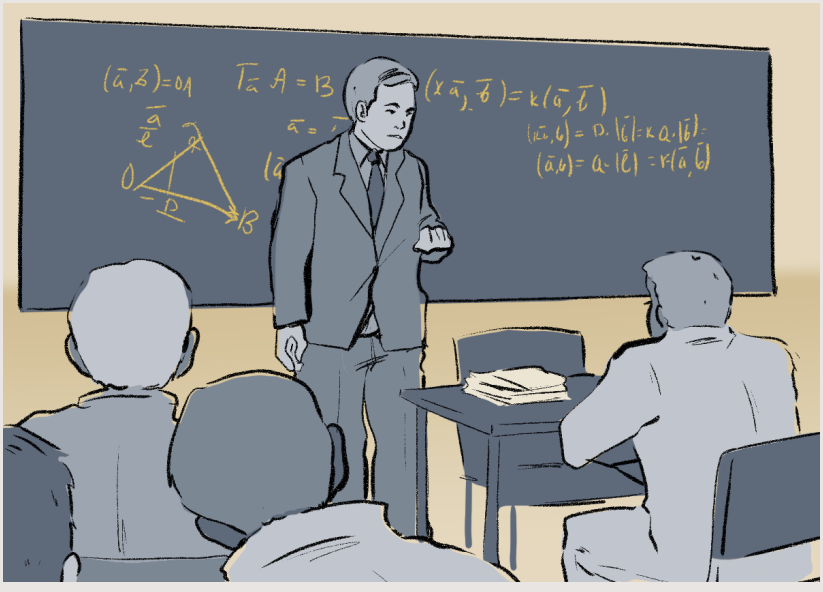

colegio-bourbaki.com